LEVEL

AD A077990

THE TRAINING OF DECISION MAKERS

Report of a Conference held at the Oregon Research Institute
8-9 July 1976
under the auspices of the
Army Research Institute for the Behavioral and Social Sciences

by

Barbara Goodman,

and

Baruch Fischhoff, Sarah Lichtenstein and Paul Slovic
Oregon Research Institute
Eugene, Oregon

2Q762722A765

DDC FILE COPY

28

March 1978

Contract DAHC 19-76-C-0028   NEW

388005

79 22 5 101

U. S. ARMY RESEARCH INSTITUTE

FOR THE BEHAVIORAL AND SOCIAL SCIENCES

A Field Operating Agency under the Jurisdiction of the
Deputy Chief of Staff for Personnel

JOSEPH ZEIDNER
Acting Technical Director

W. C. MAUS
COL, GS
Commander

Research accomplished
under contract to the Department of the Army

Oregon Research Institute

THE TRAINING OF DECISION MAKERS

Report of a Conference held at the Oregon Research Institute
8-9 July 1976
under the auspices of the
Army Research Institute for the Behavioral and Social Sciences

by

Barbara Goodman

and

Baruch Fischhoff, Sarah Lichtenstein and Paul Slovic
Oregon Research Institute
Eugene, Oregon

March 1978

Contract DAHC 19-76-C-0028

Prepared for

ari

**U.S. ARMY RESEARCH INSTITUTE
for the BEHAVIORAL and SOCIAL SCIENCES
5001 Eisenhower Avenue
Alexandria, Virginia 22333**

Army Project Number 2Q762722A765   Information Integration
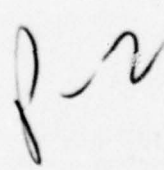
This report on an exploratory development research project is designed to meet
military management requirements for research bearing on a specific management
problem. A limited distribution is made, primarily to the operating agencies
directly involved.

ABSTRACT

The past 15 years have seen important advances in the understanding of behavioral decision processes. Despite these advances there has been surprisingly little attention paid to the training of decision makers. This report provides a summary of a conference held in July 1976 in Eugene, Oregon to evaluate the state of knowledge regarding the training of decision makers, identify important research questions and recommend directions for future research. The report provides both a brief review of current training programs, including a discussion of key attitudes and skills required for decision making, and a list of the most important research questions that emerged during the conference; Three aspects of the training problem recommended for immediate attention were (1) training specific skills, (2) evaluating the quality of decisions, and (3) implementing the knowledge obtained through decision research.

List of Participants

Berndt Brehmer, Department of Psychology, University of Umea; Sweden

Robyn M. Dawes, Oregon Research Institute, Eugene, Oregon

Hillel Einhorn, Graduate School of Business, University of Chicago

Baruch Fischhoff, Oregon Research Institute, Eugene, Oregon

Barbara Goodman, Huntington Woods, Michigan

David Gustafson, Center for Health Systems Research and Analysis,
          University of Wisconsin

Kenneth Hammond, Institute of Behavioral Science, University of Colorado

Paul J. Hoffman, Oregon Research Institute, Eugene, Oregon

Owen Jacobs, U.S. Army Research Insitute, Fort Leavenworth, Kansas

Edgar M. Johnson, U.S. Army Research Institute, Arlington, Virginia

Sarah Lichtenstein, Oregon Research Institute, Eugene, Oregon

Judith Selvidge, Decisions & Designs, Inc., McLean, Virginia

James Shanteau, Department of Psychology, Kansas State University

Lee Shulman, Department of Psychology, Michigan State University

Paul Slovic, Oregon Research Institute, Eugene, Oregon

Carl-Axel Stael von Holstein, Stanford Research Institute, Stockholm

Joseph Zeidner, U.S. Army Research Institute, Arlington, Virginia

TABLE OF CONTENTS

I. Introduction and General Issues

The rise of automation and the increased pace and intensity of contemporary conflict have made intellectual skills, especially those of judgment and decision making, crucial elements in the functioning of military and defense systems. Because of this, a considerable amount of recent research has attempted to describe the cognitive processes of decision makers and to develop techniques for improving decisions. Despite this effort, there has been surprisingly little attention paid to the training of decision makers. The present conference was convened to tackle this problem directly, with an eye towards recommending directions for future research.

One question that this conference addressed was "What do we know about training decision makers?" The question proved more complex than it first appeared to be. A great deal of time was spent in defining necessary terms, establishing appropriate boundaries, and in essence, redefining the question.

Who is a Decision Maker?

It was first necessary to define the locus of decision making. Is the decision maker a single individual or a group? Which person in a group or organization should be trained? In resolving these issues, we agreed on the following points: (1) decision making is often an interactive process within an organization; the training of a lone decision maker will be different from the training of a decision maker and his or her supporting staff; (2) when complex situations force decision makers to assign subparts of the decision task to others, the added communication and monitoring must be considered in the training program; (3) for the purposes of this conference, a decision maker is the lowest level person within an organization who could "sign a check" for/or authorize the distribution of resources. In this perspective we were concerned with tactical as opposed to strategic or operational decisions. Decision makers falling within our definition include the squad leader or military commander making tactical decisions on the basis of rapidly received, incomplete and unreliable, and sometimes irrelevant information.

General Education vs. Specific Training

Should decision makers be educated in general principles, or trained in techniques applicable to their special situations, or both? This empirical question will only be solved by experiments investigating the relative effectiveness, retention and transferability of training regimens with different balances of general principles and specific techniques.

Five levels of training, each with its respective goals were differentiated:

(1) The attitudinal level: training designed to get the decision maker to believe that better decisions arise from decomposing a problem into its relevant parts and assessing the relevant uncertainties and utilities.

(2) The aids level: the decision maker accepts the decision theoretic philosophy and learns to provide the necessary inputs into pre-constructed models.

(3) The cookbook level: teaching the decision maker to follow simple, explicit rules for making decisions. To avoid the danger of "a little knowledge being a dangerous thing," training for this level of achievement would have to include explicit warnings of the kind "If your problem has any of these elements, call your well-trained decision analyst."

(4) The formal axiomatic theory level: teaching the decision maker not only how to do a decision analysis, but the formal axiomatic theory behind it as well.

(5) The expert level: the level of training needed to practice the art to the limit of its present state of development.

Task Parameters

Feedback. Before deciding on the curriculum for training a specific decision maker, several aspects of his or her task must be taken into account. One aspect is the kind of feedback available. Two important feedback dimensions are validity and time lag, which taken together, determine feedback effectiveness.

Performance criteria. A crucial aspect of any training program is the ability to evaluate its success. For prediction tasks where people

make forecasts about future events (or past events whose outcomes are unknown to them), a good decision is indicated either by a single correct forecast or by a set of well-calibrated forecasts. Given such empirically assessable criteria, evaluating training is relatively easy with a straight-forward before-after experimental design. If subjects score significantly better after training, than we can conclude that they have learned some-thing about this task for the moment. Whether "something" is also "enough" is a separate question. To determine whether the training conferred a transferable skill, pre- and post-training measures on different tasks in different contexts are needed.

The real problem is evaluating training in situations where there is no measure of the goodness of an outcome. One method is to construct analogous laboratory tasks (e.g., simulations) which do have evaluative criteria, and then generalize the results of testing in this environment to the real world. The critical question with this approach is "Have you built sufficient subtlety and complexity into the laboratory task such that you have faith in the generalizability of the results?"

A second strategy is to use expert judges or peer ratings to determine quality. There are three kinds of peer-review approaches, the process-oriented, the problem-oriented and the outcome-oriented. With the process-oriented approach, the evaluator establishes procedural rules such that decision makers following these rules are considered good decision makers. The problem-oriented approach tries to identify and correct recurrent problems. The outcome-oriented approach concentrates on evaluating the outcomes of the decision maker's previous decisions. There are problems with each of these approaches.

The major flaw in the process- and the problem-oriented approaches is that the evidence of improved decision making is indirect. One cannot answer the question, "How good (or how much better) was performance?", following a program, in order to determine its cost-effectiveness. The major drawback with the outcome-oriented approach arises with probabilistic outcomes: good outcomes can result from bad decisions and bad outcomes from good decisions. Thus, unless a large sample of behavior is available, one can't always judge the quality of choices solely from the results produced.

As a result of the above considerations, the conference question, "What do we know about training decision makers?" was redefined as "What is the present state of knowledge about training decision makers and their staff to make more effective and/or efficient tactical decisions when the decision makers may not have access to the outcomes of previous decisions?"

II. The Current State of Knowledge

Current Training Programs

Although the training programs described here all have the general aim of aiding of decison makers, they differ in length of the training offered, specific material covered, and depth of coverage. None has been tested for relative or absolute effectiveness, nor for the retention or the transferability of its training.

An Idiot's Guide to Decision Making. An approach to teaching multi-attribute utility analysis to top-level executives and middle-level managers has been developed by Hillel Einhorn and Robin Hogarth. The Guide provides a set of simple rules with which the users can perform their own multi-attribute utility analyses and a set of warning signals which indicate when the rules should not be used because the model underlying them is violated. The Guide is taught in an hour and a half seminar; no formal theory or rationale for the rules is given. The program is sold by means of multiattribute examples selected to shake the decision maker's confidence in his or her clinical judgment and to produce a perceived need for training. The examples are carefully selected for relevance and interest. For example, one concerns the utility of individual professional basket-ball players.

The "Idiot's Guide" recommends the following steps:

1. List all the important attributes of the options you are considering. Try not to add attributes highly correlated with others already chosen.

2. Define all the attributes so that they are scaled monotonically in a positive direction with respect to total worth, making certain that all attributes are conditionally monotonic.

3. Weight all attributes equally (if you are very unhappy with this, use a weighting system based on a rank order of importance).

4. Form a matrix with the alternatives in the rows and the attributes in the columns. Fill the cell entries with ratings (which in some cases might be hard numbers like dollars).

5. Calculate the mean and standard deviation of each attribute (column).

6. Standardize each cell entry (subtract the column mean and divide by the standard deviation).

7. Multiply each standardized cell entry by the weight assigned to the attribute.

8. Sum the weighted standardized scores for each alternative.

9. Select either the alternative with the highest score or that alternative with the highest benefit to cost ratio.

Simplification of decision tasks into rules such as those in the "Idiots Guide" rests upon an elaborate theoretical rationale. For example, Einhorn and Hogarth (1975) have shown the effectiveness of equal weights in additive models.

DDI's OPINT. One of the techniques that Decision and Designs, Inc. (DDI) uses for training people in decision theoretic concepts is practice in conducting simplified decision analyses. The approach involves an interactive computer program that simplifies a decision analysis by focusing on a limited number of alternatives and the major causes of uncertainty (Selvidge, 1976). Another simplifying assumption of this process is that the probabilities are independent of actions. Decision makers start out by identifying the few alternatives and the one particular uncertain event that they will be considering. Trainees construct a value matrix similar to step four of Einhorn's procedure, except that cell entries are regrets. These are elicited by having the decision maker decide first on the best option. All other options are rated in comparison to this option. Once this has been done for every attribute, the attributes are weighted, not necessarily equally, and then combined additively. The computer then performs the necessary calculations and displays the expected value of each alternative. Sensitivity analyses are also demonstrated to show the implications of changed assessments.

Special Judgment Training. Hammond's (1971) approach, via social judgment theory, assumes that if people are taught the theory behind judgment analysis and then trained in increasingly difficult applications of task situations, they will eventually be able to analyze any problem properly. In practice sessions, trainees work with a computer terminal display initially presenting a simple task, involving, say, two cues. They are asked to use a linear rule to combine these cues. As training, with model feedback, proceeds, the tasks are complicated by using different weighting schemes, more variables, and different functional relationships. After several days of such training, people can do well on very complicated tasks. In the training up to this point, the variables are all specified. In the next step of training, people learn to uncover the structure of variables in different kinds of tasks. Here, the tasks are presented in written form and the trainees must extract the variables and determine the structure of these increasingly less well-specified problems. The final phase of training consists of having people evaluate the task structure of general statements found in print such as a policy statement written by Kissinger. Related procedures enable opposing sides in a conflict situation to carefully specify their policies and identify the exact sources of disagreement. In several field studies, this approach derived from social judgment theory (Hammond et. al., 1975, 1976) has successfully reduced conflict by showing it to be less than imagined originally (Hammond, 1971; Hammond & Boyle, 1970; Hammond & Brehmer, 1973).

Training in Decision Analysis. The Decision Analysis Group at Stanford Research Institute (SRI) conducts several different kinds of training programs for executives and staff. There are one-day, three-day, and five-day public seminars or private courses. The private courses differ from the open seminars in that they can be tailored to specific organizational needs. The one-day program attempts to convey an understanding of what decision analysis is and how it can be applied. Topics covered include: introductory decision analysis, problem structuring, probability encoding, profit lotteries and decision trees, information valuation, sensitivity analysis, and risk-preference encoding. Applications and examples are a primary vehicle for describing the concepts. SRI attempts to motivate participants by having them work through examples in which the majority of unaided decision makers make bad judgments. SRI has found that many decision makers

come to these seminars because they're worried about uncertainty and don't know how to handle it. Therefore, they are usually receptive to attempts to help them think logically about uncertainty.

In the three- and five-day seminars more time is spent on all topics and the decision makers actively work through different problems. About half a day is spent talking about biases in uncertainty judgments and in teaching participants how to overcome them. Interactive computer programs allowing the exploration of decision problems without the need for complex manual computations are demonstrated. In these seminars, people are expected to learn that the decision theoretic methodology exists, that uncertainties and utilities can be quantitatively estimated, and that they can begin to structure and work through their own decision problems. Three-day seminars are attended by high-level decision makers. Since middle-level managers and staff people are responsible for actually performing analyses, they typically attend the five-day seminars which provides them with the basic tools needed to begin the apprenticeship state of learning by doing.

MSU's Physician Training Program. Michigan State University's (MSU) Medical School approach to the training of physicians as described by Lee Shulman is based on a total curriculum design wherein decision making tools are integrated into the study of specific content areas. The approach assumes that the diagnostic phase of medicine consists of generating hypotheses about what the medical problem might be, distinguishing the relevant from the irrevalent features of the case, and then systematically gathering information to test and compare alternative hypotheses. The training curriculum based on this conception relies heavily on the case-study method. Students are presented with cases from each field of medicine selected either for frequency of occurrence or severity of consequences. Students are taught a system of medical documentation, called problem-oriented medical recording. This method preserves the process by which the problem was worked up and analyzed, and becomes itself a teaching aid. Students audit each others' records, thereby gaining additional experience by making their own decisions and critiquing those made by others. Thus, openness to analysis and criticism by one's peers is taught along with subject matter and information gathering skills. The specific form of the medical record is designed as a decision aid to structure thinking.

The physician is provided with highly usable permanent records of his or her experience for future reference.

Two questions concerning training were discussed relative to MSU's program. First, should decision training be context and/or setting specific? Second, would formal training in decision analysis per se be useful, except in altering decision makers to use ways for distinguishing among the typical phases of decisions and highlighting those for which decision algorithms can be most useful?

Training Livestock Judges. The Animal Sciences Department at Kansas State University runs a four year training program for livestock judges. As reported by James Shanteau, this program is apparently extraordinarily effective. For example, participants are able to consider actively a larger number of judgmental variables than has been reported elsewhere (at least 10 to 11 dimensions). Among its distinctive features: length --four years; social structure--trainees are organized into teams with coaches and competitions; feedback--participants transcribe their judgements and the reasons leading to them in great detail and then expose their considerations to peer and expert criticism; and criterion--experts' judgments. It is not clear which of these elements is crucial to their apparent prowess or how this "life experiment" can be generalized for other training programs.

Training expert decision analysts. The conference group agreed that expert analysts can be trained only by a combination of formal, scholastic study and practical experience under the tutelage of an expert practitioner. It was not clear whether the strongly felt need for guided practice is because there are skills involved that cannot be taught in the classroom, or because the forefront of decision analysis is changing too fast for testbooks and training programs to stay current.

Training Key Attitudes and Skills

Generating a list of attitudes and skills needed for decision making, for which the training research literature might be reviewed, turned out to be a rather difficult task. The following list of attitudes and skills summarizes the results of our efforts:

1. Recognizing that a decision problem exists and that it needs to be solved.

2. Legitimizing the problem; that is, getting others to recognize that the problem exists and requires resolution.

3. Defining the level, scope, and boundaries of a specific problem.

4. Identifying the problem variables, capturing all essentials, and eliminating redundancies.

5. Generating alternatives.

6. Defining criteria, goals or value dimensions that should be included in the final structure.

7. Selecting the structural model by which individual judgments will be combined into a composite judgment.

8. Acquiring information.

9. Dealing with uncertainty. This includes accepting uncertainty, encoding probabilities and probability distributions, and avoiding known biases (such as ignoring base rates).

10. Measuring utility.

11. Combining single judgments into composite judgments.

12. Learning from experience; that is, establishing organizational schemes for monitoring, evaluating, and improving one's decisions.

13. Knowing when and how to use heuristics, short cuts and decision aids.

14. Knowing one's own limitations and when to use normative experts.

15. Keeping focused on the problem to be solved.

16. Communicating.

17. Implementing.

The training of many of these skills has yet to be studied. We will describe the state of knowledge on training for each category separately, based on research results and on our own experience.

Recognizing a Decision Problem. Disagreement exists about whether *recognizing* problems is a general or context-specific skill. Some research in the use of group process techniques (e.g., nominal groups, Delphi, and brainstorming) shows that people can be trained to be better problem identifiers. However, this research is confined to situations where subjects' sole task is isolating problems (Utterbach, 1974; van de Ven & Delbecq, 1971).

We have heard that there is research that uses Sherlock Holmes stories to train people for problem sensitivity. For example, one point illustrated by Holmes is to be as alert to things that don't happen as to things that do. Another research area is the Bayesian research on people's perception of nonstationarities. Chinnis and Peterson (1968) showed that people can detect change, but their subjects were specifically made aware of the possibility of change. Consequently, their research tells us nothing about people's skill at detecting unanticipated problems. One suggested way to discover problems is to be alert for inconsistencies or disagreements among experts. For example, when the National Board of Medical Examiners found that different experts' evaluations of a physician's performance correlated only about .20, they realized that their evaluation procedure had to be changed. Thus, training people to keep records of judgments and review them for discrepancies might increase their ability to recognize a decision problem.

Legitimizing the problem. We had little to say about this topic. Several references on the importance of the legitimacy of problem solving in an organizational context are Delbecq (1975), Hage and Aiken (1970), and Shukla (1976).

Defining the Problem. The consensus of the group was that these skills only come with experience. Experimenters at Michigan State University (Allal, 1973; Elstein, Shulman & Sprafka, 1977) conducted a study showing that the speed of acquiring this skill may be enhanced by having the decision maker practice defining problems previously defined by experts. Their trainees examined both the expert's final structuring of the problems and transcribed thought processes. No test of generalizability or transferability was included in the experiment.

Generating Problem Variables. In well-structured situations, people apparently can learn to determine the relevant variables in a problem, although it is unclear whether they are learning a general or task-specific skill.

The generalizability of this skill to unstructured situations is also unclear. We do suggest, however, that techniques like the Kelly grid, factor analysis, and group problem solving will help in generating the variables or alternatives for a particular problem. A rule of thumb that

seems relevant is that when we don't know how to train people in a partic-
ular skill, we should give them practical experience in the skill, in hopes
that there is some unarticulated ability that will be acquired through
mere practice.

Defining the Criteria, Goals, or Value Dimensions. When the decision
is not well structured, that is, when all the relevant aspects are not
explicitly specified, difficulties arise. Foremost among these is the
neglect of one or more crucial factors whose relevance only becomes apparent,
sadly, after the decision has been made. An example of this is provided
by Birkin and Form (1973), who examined the after-effects of the "Zero
Defects" program. This program, adopted by more than 12,000 industrial
firms, attempted to attack the problem of defective workmanship by moti-
vating employees to do the job right the first time. The program was
based upon the following rationale: "Because of the complexity of today's
products and because of the drastic consequences of product failure, manage-
ment should use all means possible to meet customers' specification. Human
error on the job is not inevitable and employees, if properly motivated,
could maintain a desire to get a job done right the first time." Once
the program was implemented, many firms discovered they could not live
with the consequences of making quality a primary goal. As quality rose,
productivity declined, production deadlines were missed and amounts of
spoiled and scrapped goods increased. A high percentage of firms dropped
the program.

Although there are a number of techniques for specifying criteria,
goals, and value attributes, we know little about the implications of these
techniques for training. For example, Jerry Nadler (1972) has developed
a widely used technique called function expansion which is supposed to
produce a set of mutually exclusive and exhaustive goals. Despite its
popularity, it has apparently never been tested or compared with other
processes. A suggested technique for getting people to reduce the number
of criteria they want to use is forcing them to consider carefully the
question "What will you really want to know ten years from now?"

Selecting the Structural Model. Some component attitudes and skills
are:

1. Believing that decomposition results in better judgments and decisions.

2. Knowing the appropriate level of decomposition to employ.

3. Characterizing, by flow charts or some other technique, a conceptual picture of the problem.

4. Knowing when to use and how to construct simple models.

5. Knowing when and how to use more complex models.

6. Knowing how to draw and understand decision trees and matrices.

7. Understanding the differences between normative and descriptive models.

Training for these skills could range from an hour and a half's "Idiot's Guide" approach to many years of theoretic, philosophic and practical study. We believe that these skills can be learned. We don't know how much training is needed before a decision maker can handle tasks like modeling the components of an intricate hierarchical structure or discovering conditionality, nonindependence, and redundancy in data. We do know that the visible experts in the field have Ph.D.'s

Acquiring Information. The typical approach to information collection is gathering a mound of data in the hope that you'll have what you'll need. In Bayesian information purchase experiments, subjects frequently do not buy the correct amount of the right kind of information. In studies comparing experienced and inexperienced physicians, the latter overcollect irrelevant information. Other apparent biases are overvaluation of redundant information, preference for concrete over abstract information regardless of diagnosticity, failure to seek negative information, and underestimation of the diagnostic value of information whose causal implications are not obvious.

Several untested techniques were proposed for training proper information acquisition. One is to provide the cost of information purchased. A second is to require people to build a model of the process for which they're collecting the information and use it as a normative standard for data collection. Decision makers using Bayesian models, for example, would be required to determine likelihood ratios as well as information costs and then perform the appropriate calculations. A third possible

technique is having decision makers list and rank the questions for which they would have liked to have had answers before they made their last major policy decision. Acquiring information from another person is another skill that may be trainable, as is recognizing the effects of intercorrelations between environmental variables.

Dealing with Uncertainty. We agreed that the most important ideas and attitudes in this, the most heavily researched area, are (a) that the world is uncertain, (b) that uncertainty can be measured, and (c) that dealing logically with this uncertainty will make a difference in the quality of decision making.

One topic of discussion was how to train probability assessors to be better calibrated. Suggestions included training in:

1. Attending to relevant information and ignoring irrelevant information.
2. Recognizing and correcting judgmental biases.
3. Understanding the coherence rules of probability.
4. Understanding the equivalence of different response modes.
5. Using group situations in which people announce, defend and adjust their probabilities.
6. Using scoring rule feedback.
7. Using calibration curves.

We also agreed that it was important that probability assessors get a great deal of practice assessing probabilities in their particular field of expertise. The examples given should be richly detailed and as close as possible to the kinds of data that the decision makers would meet in the field. Furthermore, the examples should cover the probability continuum of low, medium and high likelihood events. Such substantive training is especially important in lieu of understanding how probability training transfers across tasks.

Measuring utility. Although the skills involved here can be quite complex, the group felt that little effort should be made to teach either the underlying theory or fancy applications. Since, for the most part, additive linear utility functions work quite well, people should be taught enough to know how to apply the linear model correctly and what special circumstances signal the need for expert assistance.

It appears that one way that decision makers avoid difficulties in situations with very non-linear utilities is with the "plausible range-feasible alternative" approach. That is, they consider as feasible alternatives only those that lie within critical threshold values. Within this range of plausible values, utility functions may often be reasonably approximated by a linear utility function. Other alternatives are not usually even analyzed. A danger with this approach is that decision makers may be inappropriately eliminating highly desirable alternatives. To reduce this danger, we should train decision makers to identify problem areas where this is likely to happen and teach them to deal with these special cases by using non-linear utility functions or appropriate heuristics.

Combining Single Judgments into a Composite Judgment. To date, neither judgment theorists nor decision analysts have answered the crucial question in this area: namely, when is it best to assess components, when is it best to make global judgments, and when is it best to iterate between these two procedures?

Learning from Experience. One thing that must be taught is basic inductive skills. The rules of scientific research advocated by philosophers of science were suggested as one source of insight into how to train inductive skills.

Another direct approach to teaching such skills is to alert people to avoid the known biases that hinder learning from past experience, including hindsight bias and the tendency to remember one's successes and forget one's failures. Simply informing people about biases, however, has not proven successful. What is needed are techniques that prevent perception and memory distortions. An enthusiastically received suggestion made at this meeting is to have people set up record-keeping systems that carefully document the considerations preceding their own decisions and their second guesses of other people's decisions. Such a record could also function as a mechanized data file from which actuarial models might eventually be developed. One potential problem with this system, or any other system involving feedback, is that feedback can be inaccurate, misleading or deceptive. As mentioned above, good decisions and good outcomes are not invariably associated with one another. Thus, training in setting up feedback systems would have to include rules for recognizing the limitations of different types of feedback. One example of elaborate record keeping is football coaching records where every play is analyzed

and its consequences recorded. Systematic study of the effectiveness of record keeping is needed.

Existing research suggests that mere instruction is insufficient for the learning of cognitive strategies. One also needs lots of practice in varied contexts over time. Each cognitive task may have an optimum learning program including lectures and book-learning, highly structured laboratory exercises with accompanying feedback, and guided, but less unstructured, clinical experience. The optimum mix of these components is an open question.

Knowing When and How to Properly Use Heuristics, Shortcuts, Decision Aids and Normative Experts. As research evidence is scanty here, practice appears to be the only tested means of learning these skills. Some interesting rules of thumb: it was reported that Ron Howard has recommended that an organization should devote 1% of the resources of a project to decision analysis. (If applied to the U.S. Gross National Product, this rule would provide over $10 billion in income for decision analysts.) Some educators believe that the probability that a decision aid or model will be used is directly proportional to its ease of use. There are many characterizable situations where decision making should be left to an algebraic model, or discrimination network-type model, or simulation Such situations usually involve repeatable tasks with a narrow focus.

Some of the issues involving when to use each of the various types of decision aids are discussed by Slovic (1977). Automated decision-analytic procedures create another whole category of training issues, particularly getting people to accept the procedure and use its results. Acceptance may be particularly problematic when the advice of the automated procedure is in some way counterintuitive.

Communication, Monitoring, and Implementation. It is one thing to train decision makers to make better decisions when they have total control over their analyses. The situation changes when the decision maker is dependent on others to gather information, generate alternatives or accomplish some other aspect of the decision analysis. Crucial skills, then, are communication, monitoring, and implementation. Organizational and social psychologists and managerial specialists are the appropriate sources of this knowledge (e.g., Shultz and Slevin, 1975; Ludke, 1976; Dalton, 1969; and Utterback, 1974).

III. Important Research Questions

The most important questions for future research that evolved during the two days of discussion were:

1. How do you define a "good" decision when you don't have an objective criterion?

2. How do you measure the effectiveness of a decision maker? How do you know when you have improved the effectiveness or efficiency of a decision maker through training?

3. Which combinations of training general principles and specific techniques are best to bring decision makers to different levels of competency?

4. Can we develop a diagnostic (interviewing) technique for surveying a particular decision-making context to find out where training should be concentrated?

5. How do the various different training programs used today differ? What are the virtues, difficulties, and limitations of each? What is the effectiveness, retention, and transferability of each of the skills learned in each of these programs?

6. When does the use of the decision analytic methodology improve clinical judgment?

7. How much apprenticeship is required to train a skilled decision analyst, in terms of time, number of supervised cases, and varieties of case experience?

8. Would it be worthwhile (in a cost-benefit sense) to implement the problem-oriented medical record procedure in other areas, such as the military?

9. How do you traing people to detect the dimensions of any unstructured situation? Is training for comprehensiveness detrimental to training for focussing on the important variables?

10. How do you train people to evaluate the finality or the completeness of their search for the dimensions or the variables of a problem?

11. When should decision aids, such as interactive computer programs, substitute for the more intensive personal training of decision makers? Is such training cost effective? When is training the decision maker in the theory of decision making cost effective?

12. How do you train decision makers to make proper tradeoffs between model sophistication and simplification, keeping in mind the widely-held belief that model use is directly proportional to ease of use?

13. How should we train decision makers to acquire information?

14. In probability assessment, what is the appropriate role of the decision maker and the decision analyst? Similarly, how different are the estimates produced by decision makers and the analysts? To what extent do different analysts differ?

15. When is it best to assess utility components, and when best to obtain global judgments?

16. Psychologists have uncovered many biases which people have that can hinder their ability to make decisions. What techniques can be used to overcome these biases?

17. What are the essential features of a record-keeping system that would document a decision analysis and its eventual outcome?

18. What are guidelines for the appropriate use of heuristics, short cuts, decision aids and normative experts?

IV. Recommendations

Given the current state of our knowledge, three aspects of the training problem demand immediate attention: Training specific skills, evaluating the quality of decisions, and implementing the knowledge obtained through decision research. Although these three areas are substantially inter-dependent, the research strategies each entails are sufficiently different to merit distinguishing between them.

Training Specific Skills

Tasks included here are: identifying judgmental biases, character-izing known biases in terms of underlying cognitive processes, determining which biases are amenable to training and which can only be compensated for mechanically, developing debiasing or bias-avoiding procedures where applicable, and developing guidelines for characterizing decision situations.

The research skills required here are those of experimental psychology. Although the problems are of an applied nature, it appears to us that they will be most parsimoniously attacked by looking for common elements in the

decision-making strategies used in different tasks. These are, we believe, basic cognitive processes. When a debiasing procedure proves successful, it means that we have identified the bias, understood its workings and communicated that understanding to the decision maker in a useful form. When it proves unsuccessful, it is hard to know where we went wrong. The failure of the relatively few attempts which have been made to reduce biases by simply telling people about them suggest that the problems may be more complex than we have realized. The sort of cognitive engineering we are concerned with will require considerable interaction with basic cognitive research, digging deeper in the mind when more straightforward debiasing procedures fail.

Such training studies as are conducted should be tightly controlled to maximize the chances of our being able to tell what works and what doesn't work in this difficult-to-research area. For similar reasons, we advocate concentrating on the training of specific skills rather than the development of full programs. The latter are typically more expensive and context-specific, and offer less opportunity for clearly assessing what are the effective and ineffective elements.

Evaluating the Quality of Decisions

A recurrent theme of the conference was that we do not know how to evaluate most important decisions. Without such knowledge, there is no way of assessing the value of the various training programs now offered or the validity of the claims made for them. At the extreme, we do not even know if any simple decision-making procedures, however flimsy its axiomatic basis, is not as good as the most sophisticated.

Developing evaluative criteria is a difficult, long-term project. It should be accompanied by an attempt to survey and compare existing programs using the best tools currently available. Such a survey should include commercially-marketed nonaxiomatic approaches like Jerry Nadler's and Dale Carnegie's. We should not exclude the possibility that they know something that we don't.

Another aspect of the evaluation problem is to assess the importance of the judgmental biases that have been identified by various investigators,

that is, to determine how much of a difference they make in the optimality of decisions. What is called for here is some general work on the sensitivity of decisions to bias. It should consider the apparently untouched question of how bias and error in the various component judgments used in a complex analysis combine to affect the final outcome. Some analysis of the ecological validity of judgmental heuristics would also be useful in order to establish the frequency and extent of the biases they induce.

A final aspect of the evaluation problem is developing ways to help people best assess the quality of their own decisions (and those of their colleagues), and learn from their own experience. Hindsight research (Fischhoff, 1975) suggests some ways in which this might be done, as do the decision record-keeping systems described by Shulman and Shanteau.

Implementing Present Knowledge

There is, already, much knowledge that can be used to train decision makers and improve the quality of their performance. For repeatable tasks, the heretofore covert decision processes of the expert can be modeled and made explicit in a way that should be quite useful to trainees. In some cases, these models will take the form of algebraic equations. In others, more complex models, such as sequentially branching computer programs, will be necessary. The potential of judgment modeling for facilitating military and defense decisions is unlimited. Demonstration projects should begin immediately. Relevant background information is provided by Hammond et al. (1975) and Slovic and Lichtenstein (1971).

The majority of important decisions will be made in unique situations that permit little time for deliberation. The standard method of decisions in these situations will continue to be intuitive. Given the pitfalls to which unaided intuitions are susceptible, we have little reason to be comfortable with this prospect. Since we cannot avoid heavy reliance on intuition, we should at least educate decision makers to the pitfalls that await the unwary. We recommend that consideration be given to the development of a curriculum designed to educate the intuitions of decision makers in the subtleties of dealing with a complex and uncertain environment.

References

Allal, L. K. Training of medical students in a problem solving skill: The generation of diagnostic problem formulations, Ph.D. Dissertation, Michigan State University, 1973.

Birkin, S. J. & Ford, J. S. The quantity/quality dilemma: The impact of a zero defects program. In J. L. Cochrane & M. Zeleny, (Eds.), Multiple Criteria Decision Making, Columbia, S. C.: Univ. of S. Carolina Press, 1973, 517-529.

Chinnis, J. O., Jr. & Peterson, C. R. Inference about a nonstationary process. Journal of Experimental Psychology, 1968, 77, 620-625.

Dalton, G. Influence and organizational change. Paper presented at the Conference on Organization Behavior Models, Kent State, University, 1969.

Delbecq, A. Negotiation mandates which increase acceptance of evaluation findings concerning demonstration programs. Proceedings of Academy of Management Meetings, August, 1975.

Einhorn, H., & Hogarth, R. Unit weighting schemes for decision making. Organizational Behavior and Human Performance, 1975, 13, 171-192.

Elstein, A. S., Shulman, L. S. & Sprafka, S. A. The Psychology of Medical Diagnosis, Cambridge, Ma.: Harvard Univ. Press, in press.

Fischhoff, B. Hindsight: Thinking backwards? Psychology Today, 1975, 8 (4), 70-76.

Hage, G. & Aiken, M. Social Change in Complex Organizations, New York: Random House, 1970.

Hammond, K. R. Computer graphics as an aid to learning. Science, 1971, 172, 903-908.

Hammond, K. R. & Boyle, P. J. R. Quasi-rationality, quarrels and new conceptions of feedback. Bulletin of the British Psychological Society, 1970, 24, 103-113.

Hammond, K. R. & Brehmer, B. Quasi-rationality and distrust: Implications for international conflict. In L. Rappoport and D. A. Summers, (Eds.), Human Judgment and Social Interaction. New York: Holt, 1973.

Hammond, K. R., Rohrbaugh, J., Mumpower, J. L. & Adelman, L. Social judgment theory: Applications in policy formation. In M. F. Kaplan & S. Schwartz, (Eds.), Human Judgment and Decision Processes: Applications in Problem Settings. New York: Academic Press, 1976.

Hammond, K. R., Stewart, T. R., Brehmer, B. & Steinmann, D. O. Social judgment theory. In M. F. Kaplan and S. Schwartz, (Eds.), Human Judgment and Decision Processes: Formal and mathematical approaches. New York: Academic Press, 1975.

Ludke, R. L. An analysis of referral decisions that clincial cancer provides. Ph.D. Dissertation proposal, Dept. of Industrial Engineering, Univ. of Wisconsin, June, 1976.

Nadler, G. Work Design. Homewood, Ill.: Irwin Press, 1972.

Selvidge, J. Rapid screening of decision options. Technical Report, TR 76-12, McLean, Va.: Decisions & Designs, Inc., 1976.

Shukla, R. Organizational influence in innovative decision. Center for Health Systems Research and Analysis Working Paper, June, 1976.

Schultz, R. & Selvin, D. Implementing management science and operations research, New York: American Elsevier, 1975.

Slovic, P. Towards understanding and improving decisions. In Science Prospectives, Volume Prepared for the 30th Anniversary of the Office of Naval Research, in press.

Slovic, P. & Lichtenstein, S. Comparison of Bayesian and regression approaches to the study of information processing in judgment. Organizational Behavior and Human Performance, 1971, 6, 649-744.

Utterback, J. M. Innovation in industry and the diffusion of technology. Science, 1974, 183:4125, 620-626.

Van de Ven, A. & Delbecq, A. Nominal vs. interacting group processes for committee decision making effectiveness. Academy of Management Journal, 1971, 14, 3.